

10/19/00

JC904 U.S. PTO

10-20-00

EK 53821679505

THE ASSISTANT COMMISSIONER FOR PATENTS  
Washington, D.C. 20231

DOCKET NUMBER: AUS920000620US1  
DATE: October 19, 2000

JC759 U.S. PTO  
09/692342  
10/19/00

Sir:

Transmitted herewith for filing is the Patent Application of:

Inventors: G. M. Nordstrom et al  
For: METHOD AND SYSTEM FOR INFORMING AN OPERATING SYSTEM IN A SYSTEM AREA NETWORK  
WHEN A NEW DEVICE IS CONNECTED

Enclosed are:

- ☒ Patent Specification and Declaration (3 Declarations)
- ☒ sheets of drawing(s). (Informal) 8
- ☒ An assignment of the invention to International Business Machines Corporation (includes Recordation Form Cover Sheet).
- ☐ A certified copy of a \_\_\_\_\_ application.
- ☐ Information Disclosure Statement, PTO 1449 and copies of references.

The filing fee has been calculated as shown below:

For	Number Filed	Number Extra	Rate	Fee
Basic Fee				\$710.00
Total Claims	22 - 20	2	x 18 =	\$ 36.00
Indep. Claims	4- 3	1	x 80 =	\$ 80.00
MULTIPLE DEPENDENT CLAIM PRESENTED x 270 =				\$ 0.00
TOTAL				\$826.00

- ☒ Please charge my Deposit Account No. 09-0447 in the amount of \$ 826.00 . A duplicate copy of this sheet is enclosed.
- ☒ The Commissioner is hereby authorized to charge payment of the following fees associated with this communication or credit any overpayment to Deposit Account 09-0447 . A duplicate copy of this sheet is enclosed.
- ☒ Any additional filing fees required under 37 CFR 1.16.
- ☒ Any patent application processing fees under 37 CFR 1.17.

Respectfully submitted,

By Mark E. McBurney  
Mark E. McBurney  
Registration No. 33,114  
Intellectual Property Law Dept.  
IBM Corporation  
11400 Burnet Road, Zip 4054  
Austin, Texas 78758  
Telephone (512) 823-1003

**METHOD AND SYSTEM FOR INFORMING AN OPERATING SYSTEM IN A  
SYSTEM AREA NETWORK WHEN A NEW DEVICE IS CONNECTED**

**BACKGROUND OF THE INVENTION**

The present invention is related to the subject matter of the following commonly assigned, co-pending United States Patent Applications filed concurrently herewith: Serial No. \_\_\_\_\_ (Docket No. AUS9-2000-0621) entitled "Efficient Process for Handover Between Subnet Managers"; Serial No. \_\_\_\_\_ (Docket No. AUS9-2000-0622) entitled "Method and System for Scalably Selecting Unique Transaction Identifiers"; Serial No. \_\_\_\_\_ (Docket No. AUS9-2000-0623) entitled "Method And System For Reliably Defining and Determining Timeout Values In Unreliable Datagrams"; and Serial No. \_\_\_\_\_ (Docket No. AUS9-2000-0624) entitled "Method and System For Choosing A Queue Protection Key That is Tamper-proof From An Application". The content of the above-referenced applications is incorporated herein by reference.

**1. Technical Field:**

The present invention relates in general to computer networks and, in particular, to connecting devices to a computer network. Still more particularly, the present invention relates to a method and system for informing an operating system when a new device is connected to a System Area Network.

**2. Description of the Related Art:**

Distributed computer networks are known in the art. In a traditional network, individual components of the

network are interconnected via a parallel bus, such as a PCIX bus. The parallel bus has a relatively small number of plug-in ports for connecting the components. The number of plug-in-ports is set (i.e., the number cannot be increased). At maximum loading, a PCIX bus transmits data at about 1 Gbyte/second.

The introduction of high performance adapters (e.g., SCSI adapters), Internet-based networks, and other high performance network components has resulted in increased demand for bandwidth, faster network connections, distributed processing functionality, and scaling with processor performance. These and other demands are quickly outpacing the current parallel bus technology and are making the limitations of parallel buses even more visible. PCIX bus, for example, is not scalable, i.e., the length of the bus and number of slots available at a given frequency cannot be expanded to meet the needs for more components, and the limitation hinders further development of fast, efficient distributed networks, such as system area networks. New switched network topologies and systems are required to keep up with the increasing demands.

The present invention recognizes the need for faster, more efficient computer networks offering the features demanded by the developments of technology. More specifically, the present invention recognizes the need for providing a method or system for informing an operating system (OS) when a (new) device is connected to a scalable, distributed computer network.

### SUMMARY OF THE INVENTION

A method and system are provided for dynamically informing an OS of a distributed computer system, when a (new) device is added on the network. The invention is thus applicable to a distributed computing system, such as a system area network, having end nodes, switches, and routers, and links interconnecting these components. The switches and routers interconnect the end nodes and route packets, i.e., sub-components of messages being transmitted, from a source end node to a target end node. The target end node then reassembles the packets into the message.

In a preferred embodiment of the invention, an OS registers with the subnet administrator (SA) to be notified when a new component is added to the subnet and placed in a partition to which the OS has access. The subnet manager (SM) initializes and configures the new component that is added and, through the partition manager, associates the component to one or more partitions of the subnet. The SA then notifies the OSs with access to the partition(s) to which the component is associated and that have a right to access the new component that the new component is available to the OS. Following, the OS updates a database of available components to include the new component and establishes communication with the new component.

In one embodiment, the SM periodically sweeps the network to search for and discover new components. In another embodiment, the new component alerts the SM when the new component is first connected to the network.

100% 99% 98% 97% 96% 95% 94% 93% 92% 91% 90% 89% 88% 87% 86% 85% 84% 83% 82% 81% 80% 79% 78% 77% 76% 75% 74% 73% 72% 71% 70% 69% 68% 67% 66% 65% 64% 63% 62% 61% 60% 59% 58% 57% 56% 55% 54% 53% 52% 51% 50% 49% 48% 47% 46% 45% 44% 43% 42% 41% 40% 39% 38% 37% 36% 35% 34% 33% 32% 31% 30% 29% 28% 27% 26% 25% 24% 23% 22% 21% 20% 19% 18% 17% 16% 15% 14% 13% 12% 11% 10% 9% 8% 7% 6% 5% 4% 3% 2% 1% 0%

## BRIEF DESCRIPTION OF THE DRAWINGS

The novel features believed characteristic of the invention are set forth in the appended claims. The invention itself however, as well as a preferred mode of use, further objects and advantages thereof, will best be understood by reference to the following detailed description of an illustrative embodiment when read in conjunction with the accompanying drawings, wherein:

**Figure 1** depicts a system area network (SAN) in which the present invention is preferably implemented;

**Figure 2** is a diagram of software aspects of an exemplary host processor end node for the SAN of **Figure 1** in accordance with the present invention;

**Figure 3** is a diagram of an exemplary host channel adapter of the SAN of **Figure 1** in accordance with the present invention;

**Figure 4** is a diagram of the work queue and completion queue process for the queue pairs (QPs) of **Figure 3**;

**Figure 5** is a diagram of the software aspects of SAN management model in accordance with the present invention;

**Figure 6** is a diagram of the messaging between the subnet manager, OS and a newly added component to effect the preferred embodiment of the present invention.

5

[illegible]

## DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

The present invention provides a discovery utility that allows an OS of a distributed computer system, such as a system area network (SAN), to be notified when ever a new component (node or device) is added to the SAN. The invention is also applicable to discovery of previously connected devices that were in the OS database but have been removed from the network for one reason or another. The following description will be understood to refer to both types of devices although specific embodiments are described with reference to a new device. In order to appreciate the environment within which the invention is preferably practiced, a description of a SAN configured with routers, switches, and end nodes, etc. is provided below. Presentation of the environment and particular functional aspects of the environment which enable the invention to be practiced are provided with reference to **Figures 1-5**. Section headings have been provided to distinguish the hardware and software architecture of the SAN. However, those skilled in the art understand that the descriptions of either architecture necessarily includes references to both components.

### SAN HARDWARE ARCHITECTURE

With reference now to the figures and in particular with reference to **Figure 1**, there is illustrated an exemplary embodiment of a distributed computer system. Distributed computer system **100** represented in **Figure 1** is provided merely for illustrative purposes, and the



embodiments of the present invention described below can be implemented on computer systems of numerous other types and configurations. For example, computer systems implementing the present invention may range from a small server with one processor and a few input/output (I/O) adapters to very large parallel supercomputer systems with hundreds or thousands of processors and thousands of I/O adapters. Furthermore, the present invention can be implemented in an infrastructure of remote computer systems connected by an Internet or intranet.

As shown in **Figure 1**, distributed computer system **100** includes a system area network (SAN) **113**, which is a high-bandwidth, low-latency network interconnecting nodes within the distributed computer system. More than one (1) SAN **113** may be included in a distributed computer system **100** and each SAN **113** may comprise multiple sub-networks (subnets).

A node is herein defined to be any component that is attached to one or more links of a network. In the illustrated distributed computer system, nodes include host processors **101**, redundant array of independent disks (RAID) subsystem **103**, I/O adapters **105**, switches **109A-109C**, and router **111**. The nodes illustrated in **Figure 1** are for illustrative purposes only, as SAN **113** can connect any number and any type of independent nodes. Any one of the nodes can function as an end node, which is herein defined to be a device that originates or finally consumes messages or frames in the distributed computer system **100**.

SAN **113** is the communications and management infrastructure supporting both I/O and inter-processor communications (IPC) within distributed computer system **100**. Distributed computer system **100**, illustrated in **Figure 1**, includes a switched communications fabric (i.e., links, switches and routers) allowing many devices to concurrently transfer data with high-bandwidth and low latency in a secure, remotely managed environment. End nodes can communicate over multiple ports and utilize multiple paths through SAN **113**. The availability of multiple ports and paths through SAN **113** can be employed for fault tolerance and increased-bandwidth data transfers.

SAN **113** includes switches **109A-109C** and routers **111**. Switch **109A-109C** connects multiple links together and allows routing of packets from one link to another link within SAN **113** using a small header Destination Local Identifier (DLID) field. Router **111** is capable of routing frames from one link in a first subnet to another link in a second subnet using a large header Destination Globally Unique Identifier (DGUID). Router **111** may be coupled via wide area network (WAN) and/or local area network (LAN) connections to other hosts or other routers.

In SAN **113**, host processor nodes **101** and I/O nodes **106** include at least one Channel Adapter (CA) to interface to SAN **113**. Host processor nodes **101** include central processing units (CPUs) **119** and memory **121**. In one embodiment, each CA is an endpoint that implements

the CA interface in sufficient detail to source or sink packets transmitted on SAN **113**. As illustrated, there are two CA types, Host CA (HCA) **117** and Target CA (TCA) **127**. HCA **117** is used by general purpose computing nodes to access SAN **113**. In one implementation, HCA **117** is implemented in hardware. In the hardware implementation of HCA **117**, HCA hardware offloads much of CPU and I/O adapter communication overhead. The hardware implementation of HCA **117** also permits multiple concurrent communications over a switched network without the traditional overhead associated with communicating protocols. Use of HCAs **117** in SAN **113** also provides the I/O and IPC consumers of distributed computer system **100** with zero processor-copy data transfers without involving the operating system kernel process. HCA **117** and other hardware of SAN **113** provide reliable, fault tolerant communications.

The I/O chassis **106** includes I/O adapter backplane and multiple I/O adapter nodes **105** that contain adapter cards. Exemplary adapter cards illustrated in **Figure 1** include SCSI adapter card **123A**, adapter card **123B** to fiber channel hub and FC-AL devices, Ethernet adapter card **123C**, graphics adapter card **123D**, and video adapter card **123E**. Any known type of adapter card can be implemented. The I/O chassis **106** also includes switch **109B** in the I/O adapter backplane to couple adapter cards **123A-123E** to SAN **113**.

RAID subsystem **103** includes a microprocessor **125**, memory **126**, a Target Channel Adapter (TCA) **127**, and

multiple redundant and/or striped storage disks **129**.

In the illustrated SAN **113**, each link **115** is a full duplex channel between any two network elements, such as end nodes, switches **109A-109C**, or routers **111**. Suitable links **115** may include, but are not limited to, copper cables, optical cables, and printed circuit copper traces on backplanes and printed circuit boards. The combination of links **115** and switches **109A-109C**, etc. operate to provide point-to-point communication between nodes of SAN **113**.

## **SAN SOFTWARE ARCHITECTURE**

### **Software Components**

Software and hardware aspects of an exemplary host processor node **101** are generally illustrated in **Figure 2**. Host processor node **101** includes one or more processors that execute a set of consumer processes **201**. Host processor node **101** includes HCA **117** with ports **205**. Each port **205** connects to a link **115** of SAN **113**. Ports **205** can connect to one SAN subnet or multiple SAN subnets. Utilizing message and data services **203**, consumer processes **201** transfer messages to SAN **113** via verbs interface **207**. Verbs interface **207** is generally implemented with an operating-system specific programming interface.

A software model of HCA **117** is illustrated in **Figure 3**. HCA **117** includes a set of queue pairs (QPs) **301**, which transfer messages across ports **205** to the subnet. A single HCA **117** may support thousands of QPs **301**. By

contrast, TCA **127** in an I/O adapter typically supports a much smaller number of QPs **301**. Also illustrated are subnet management administration (SMA) **209**, management packets **211** and a number of virtual lanes **213**, which connect transport layer with ports **205**.

Turning now to **Figure 5**, there is illustrated a software management model for nodes on SAN **113**. SAN architecture management facilities provides a Subnet Manager (SM) **303A**, a Subnet Administration (SA) **303B**, and an infrastructure that supports a number of general management services. The management infrastructure includes a Subnet Management Agent (SMA) **307** operating in each node and defines a general service interface that allows additional general services agents. Also, SAN architecture defines a common management datagram (MAD) message structure for communicating between managers and management agents.

SM **303A** is responsible for initializing, configuring and managing switches, routers, and channel adapters. The SM can be implemented within other devices, such as a channel adapter or a switch. One SM **303A** of SAN is dedicated as a master SM and is responsible for:

- discovering the subnet topology; configuring each channel adapter port with a range of Local Identification (LID) numbers, Global Identification (GID) number, subnet prefix, and Partition Keys (P\_Keys); configuring each switch with a LID, the subnet prefix, and with its forwarding database; and maintaining the end node and service databases for the subnet to provide a Global

Unique Identification (GUID) number to LID/GID resolution service as well as a services directory. Thus, management of SAN **113** and SAN components, such as HCAs **117**, TCAs (or end nodes) **127**, switches **109**, and routers **111** are completed utilizing Subnet Management (SM) **303A** and Subnet Administration (SA) **303B**. SMPs are used to discover, initialize, configure, and maintain SAN components through management agents **307** of end nodes **305**. SAN SA packets are used by SAN components to query and update subnet management data. Control of some aspects of the subnet management are provided via a user management console **311** in host-based end node **309**.

#### MESSAGE TRANSFER PROCESS

SAN **113** provides the high-bandwidth and scalability required for I/O and also supports the extremely low latency and low CPU overhead required for Interprocessor Communications (IPC). User processes can bypass the operating system (OS) kernel process and directly access network communication hardware, such as HCAs **117**, which enable efficient message passing protocols. SAN **113** is suited to current computing models and is a building block for new forms of I/O and computer cluster communication. SAN **113** allows I/O adapter nodes **105** to communicate among themselves or communicate with any or all of the processor nodes **101** in the distributed computer system. With an I/O adapter attached to SAN **113**, the resulting I/O adapter node **105** has substantially the same communication capability as any processor node **101** in the distributed computer system.

For reliable service types of messages, end nodes, such as host processor nodes **101** and I/O adapter nodes **105**, generate request packets and receive acknowledgment packets. Switches **109A-109C** and routers **111** pass packets along from the source to the target (or destination). Except for the variant CRC trailer field, which is updated at each transfer stage in the network, switches **109A-109C** pass the packets along unmodified. Routers **111** update the variant CRC trailer field and modify other fields in the header as the packet is routed.

In SAN **113**, the hardware provides a message passing mechanism that can be used for Input/Output (I/O) devices and Interprocess Communications (IPC) between general computing nodes. Consumers (i.e., processing devices connected to end nodes) access SAN **113** message passing hardware by posting send/receive messages to send/receive work queues (WQ), respectively, on a SAN Channel Adapter (CA).

A message is herein defined to be an application-defined unit of data exchange, which is a primitive unit of communication between cooperating processes. A packet (or frame) is herein defined to be one unit of data encapsulated by networking protocol headers (and trailer). The headers generally provide control and routing information for directing the packet (or frame) through SAN **113**. The trailer generally contains control and cyclic redundancy check (CRC) data for ensuring that frames are not delivered with corrupted content.

Consumers use SAN verbs to access HCA functions.

The software that interprets verbs and directly accesses the CA is known as the Channel Interface (CI) **219**. Send/Receive work queues (WQ) are assigned to a consumer as a Queue Pair (QP). Messages may be sent over five different transport types, Reliable Connected (RC), Reliable Datagram (RD), Unreliable Connected (UC), Unreliable Datagram (UD), and Raw Datagram (RawD). Consumers retrieve the results of these messages from a Completion Queue (CQ) through SAN send and receive work completions (WC). The source CA takes care of segmenting outbound messages and sending them to the destination. The destination or target CA takes care of reassembling inbound messages and placing them in the memory space designated by the destination's consumer. These features are illustrated in the figures below.

Referring now to **Figure 4**, there is illustrated a block diagram of work and completion queue processing. Each QP **301** provides an input to a Send Work Queue (SWQ) **407** and a Receive Work Queue (RWQ) **409**. SWQ **407** sends channel and memory semantic messages, and RWQ **409** receives channel semantic messages. A consumer calls a verb (within verbs interface **207**) to place Work Requests (WRs) into a WQ. A Send WR **403** is a channel semantic operation to push a set of local data segments **417** to the data segments referenced by a remote node's Receive WQE **405**. Each of the Send WR's data segments **417** contains a virtually contiguous memory region. The virtual addresses used to reference the local data segments **417** are in the address context of the process that created the local QP **301**.



As shown in **Figure 4**, WRs **403** that have been placed onto a WQ by consumer process **401** are referred to as work queue elements (WQEs) **405**. WQEs **405** are executed by hardware **415** in HCA **117**. SWQ **407** contains WQEs **405** that describe data to be transmitted on the SAN fabric. RWQ **409** contains WQEs **405** that describe where to place incoming channel semantic data received from SAN **113**.

In one embodiment, Receive Work Queues **409** only support one type of WQE **405**, which is referred to as a receive WQE. The receive WQE provides a channel semantic operation describing a local memory space into which incoming send messages are written. The receive WQE includes a scatter list describing several virtually contiguous memory spaces. An incoming send message is written to these memory spaces. The virtual addresses are in the address contexts of the process that created the local QP **301**.

Verbs interface **207** also provide a mechanism for retrieving completed work from completion queue **411**. Completion queue **411** contains Completion Queue Elements (CQEs) **413** which **413** contain information about previously completed WQEs **405**. Completion queues **413** are employed to create a single point of completion notification for multiple QPs **301**. CQE **413** contains sufficient information to determine the QP **301** and specific WQE **405** that completed. A completion queue context (not shown) is a block of information that contains pointers to, length, and other information needed to manage individual completion queues **411**.

## REMOTE OPERATION FUNCTIONALITY

SAN 113, with its interlinked arrangement of components and sub-components, provides a method for completing remote operations, by which processor nodes may directly control processes in I/O nodes. Remote operation also permits the network to manage itself. A remote direct memory access (RDMA) Read WR provides a memory semantic operation to read a virtually contiguous memory space on a remote node. A memory space can either be a portion of a memory region or a portion of a memory window. A memory region references a previously registered set of virtually contiguous memory addresses defined by a virtual address and length. A memory window references a set of virtually contiguous memory addresses which have been bound to a previously registered region.

The RDMA Read WR writes the data to a virtually contiguous local memory space. Similar to Send WR 403, virtual addresses used by the RDMA Read WQE to reference the local data segments are in the address context of the process that created the local QP 301. The remote virtual addresses are in the address context of the process owning the remote QP targeted by the RDMA Read WQE.

RDMA Write WQE provides a memory semantic operation to write a virtually contiguous memory space on a remote node. RDMA Write WQE contains a scatter list of local virtually contiguous memory spaces and the virtual address of the remote memory space into which the data from the local memory spaces is written.

RDMA FetchOp WQE provides a memory semantic operation to perform an atomic operation on a remote

word. RDMA FetchOp WQE is a combined RDMA Read, Modify, and Write operation. RDMA FetchOp WQE can support several read-modify-write operations, such as "Compare and Swap if Equal."

5

A Bind (unbind) remote access key (R\_Key) WQE provides a command to the HCA hardware to modify a memory window by associating the memory window to a memory region. A second command to destroy a memory window by disassociating the memory window to a memory region is also provided. The R\_Key is part of each RDMA access and is used to validate that the remote process has permitted access to the buffer.

10

15

#### INFORMING OPERATING SYSTEMS WHEN A NEW DEVICE IS DISCOVERED ON SAN

**Figure 6** illustrates an illustrative embodiment of hardware and software components relevant to the device discovery process of the present invention. As illustrated, SM Node **601**, Added Node **605**, and Host Node **603** are interconnected via one or more switches **109**. SM Node **601** comprises SM **607** and SA **609**. Although SA **609** is illustrated as a separate component from SM **607**, in other implementations, SA **609** may be included in SM **607**. Host Node **603** comprises consumer OS **611**.

20

25

30

The preferred embodiment of the present invention provides a dynamic signaling mechanism for informing an OS when a new device (node or component) has been added to SAN **113**. Whenever a new component, such as a storage adapter, is added to a subnet, subnet manager (SM) **607** is

utilized to configure the attributes of the new component. Configuring the component includes associating the new component with partitions within the subnet and making the new component available to one or more OSs on the same subnet. In one embodiment, the associating of the component to partitions is completed by a partition manager (or agent). As utilized herein the term partition refers to the creation of logical barriers between component groups within the subnet. An operation or component in a first partition is isolated from interaction with components or operations in a second partition unless that component or operation in the first partition has specific privilege to access the second partition.

While SAN 113 is operational, SM 607 periodically scans the network for devices that are added after the initial configuration sweeps. SAN 113 is queried and configured by means of SMPs. SM 607, by means of the SMPs, obtains information about the SAN components and sets configuration values in the components to configure SAN 113 to exhibit particular characteristics. Some information utilized to determine the SAN component and network characteristics includes assigned addresses, partitioning information, Service Level (SL) to Virtual Lane (VL) Mapping, VL arbitration table, Light Emitting Diode (LED) state information, and switch routing tables.

On SAN 113, if a new device is added, one or more OSs need to be aware of the device in order to provide support to the device. The present invention provides the mechanism by which the OS is informed about a new device added to SAN 113.

Referring now to **Figure 8**, there is illustrated an illustrative embodiment of a process of registering an OS for notification of new devices on SAN **113**. The subscription process starts at block **803** and proceeds to block **805**, where a consumer (e.g., an OS) that desires to be informed of changes to SAN **113** sends a Management Datagram (MAD) packet to SA **609** requesting to be informed of changes in SAN **113**. Following, in block **807**, SA **609** checks to see if the consumer has privileges to the paths to which the consumer is subscribing. The consumer, for example, may not have access because the paths that are being requested are owned by a partition of SAN **113** to which the consumer does not belong. Thus, in one embodiment, an OS that has signed up for notification is notified by the SA only when the new component is added to the subnet is placed in a partition to which the OS has access.

If a determination is made at block **807** that the consumer does not have access rights (or privileges), a rejection packet is sent back to the consumer at block **811** informing the consumer that the request was rejected. Then, the process terminates at block **815**. If, however, the consumer does have the correct access privileges, the process continues to block **809** where SA **609** places the subscription information in a database utilized to keep track of which consumers want to be informed of network changes. SA **609** then sends a subscription acceptance packet to the consumer at block **813**, and the operation is completed at block **815**.

In another embodiment, the determination of access privileges at block 807 is delayed until a change is detected in the configuration of SAN 113. Once the configuration change is registered by SM 607, the determination of whether or not the consumer has the privileges to be informed of the change is made. If the consumer does not have the required privilege, then the consumer is not informed of the configuration change. This latter embodiment has the advantage of allowing changes in access privileges in between the time the consumer is registered for notification and the time of the configuration change of SAN 113. Consequently, changes in access privileges of the consumer (OS' subscriptions), which occur after the consumer subscribes to be notified but before the component is added are accounted for. When the subscription operation is complete, the consumer is ready to be informed of any changes.

Continuing now to **Figure 7**, there is illustrated an illustrative embodiment of an added node discovery process. The discovery process starts at block 703 and continues to block 705, where a new node is plugged into SAN 113. The new node is detected in block 707 by one of two methods. In the first method, SM 607 discovers the added node during one of its scheduled sweeps of SAN 113. In the second method, the added node supports Trap packets and sends a Trap message to SM 607 when it is first connected. Both methods are described above.

After SM 607 discovers added node 605, SM 607 configures added node 605 in block 709 to set up the

addressing and other information about added node 605 as in the normal SAN configuration process. The setup process includes assigning port attributes (LID, service level, SL, etc.) to the added ports and assigning partition table entries to the partition tables of the subnet.

SM 607 then passes the information about added node to SA 609 at block 711. SA 609 logs the added component 605 in a database of known (or discovered) components at block 713 and then checks at block 715 to see if any consumer is subscribed to be notified that an added node has been discovered. If no consumer is subscribed, then the process completes at block 719; otherwise SA 609 sends a report packet to the subscribed consumer at block 717, providing the subscribed consumer enough information about the added node to enable the consumer to utilize (or communicate with) the added node.

In one embodiment, the check performed in decision block 715 may be replaced by a check to determine which of the consumers, if any, that subscribed to SA 609 have the correct access privileges for the partition in which added node 605 exists. Then, only consumers with the correct access privileges are informed of added node 605.

Referring again to **Figure 6**, an example of messaging (i.e., the flow of messages) among nodes on SAN 113 during operation of the invention is illustrated. Consumer OS 611 in Host node 603 desires to be notified of changes in SAN 113, for example, addition of a

component via Added Node 605, so that the consumer OS 611 may support the component and provide the component with access to SAN resources, accordingly. These changes may also include removals of components or nodes, but the operating scenario described herein focuses on the addition of components.

In the illustrated embodiment, consumer OS 611 issues a SA inform packet 615 on SAN 113. SA inform packet 615 travels via switch 109 and links 115 to SA 609. In one embodiment, when a newly added node 605 is connected to SAN 113, added node 605 has the connection to SAN 113 detected by switch 109, and, if switch 109 supports the sending of a trap on a change to the status of one of the ports, then switch 109 sends a trap packet 613 to master SM 607 of SAN 113. In a next embodiment, when switch 109 does not support Trap packets, added node 605 is detected when master SM 607 completes a sweep of SAN 113. In this later embodiment, master SM 607 sweeps SAN 113 on a periodic basis looking for changes in the configuration of SAN 113. The periodicity of the sweeps is a network parameter that is adjustable by a subnet administrator.

When master SM 607 detects a change, either through a Trap packet or on its own, master SM 607 informs SA 609 of the change. SA 609 then sends a report packet 617 to subscribed consumer OS 611. Thus, consumer OS 611 can perform the necessary setup procedures to thereafter access the added component 605. In the preferred embodiment, the OS updates a database of available



components and establishes communication with the new component.

Accordingly, with the present invention, new nodes  
5 are dynamically discovered across SAN 113 and configured  
for use by the a remote consumer (or OS) that is  
ultimately going to support and/or control the new node.  
A system administrator can thus connect a new node to any  
location of SAN 113 and the new node is automatically  
10 provided necessary software support.

As a final matter, it is important to note that  
while an illustrative embodiment of the present invention  
15 has been, and will continue to be, described in the  
context of a fully functional data processing system,  
those skilled in the art will appreciate that the  
software aspects of an illustrative embodiment of the  
present invention are capable of being distributed as a  
20 program product in a variety of forms, and that an  
illustrative embodiment of the present invention applies  
equally regardless of the particular type of signal  
bearing media used to actually carry out the  
distribution. Examples of signal bearing media include  
25 recordable type media such as floppy disks, hard disk  
drives, CD ROMs, and transmission type media such as  
digital and analogue communication links.

In the above detailed description of the preferred  
30 embodiments, reference is made to the accompanying  
drawings which form a part hereof, and in which is shown  
by way of illustration specific embodiments in which the  
invention may be practiced. It is to be understood that  
other embodiments may be utilized and structural or

5

## CLAIMS

What is claimed is:

1        1.    A method for configuring a network, said method  
2        comprising the steps of:

3  
4                dynamically determining when a component is added to  
5        a node of said network; and

6  
7                in response to said dynamically determining step,  
8        configuring said network to provide support for said  
9        component.

1        2.    The method of Claim 1, further comprising the steps  
2        of:

3                registering an OS with an management system of said  
4        network to provide a notification when said component is  
5        added; and

6                alerting said OS that said component is added during  
7        said configuring step.

1        3.    The method of Claim 2, wherein said dynamically  
2        determining step is completed by said management system  
3        and includes the step of monitoring a network via a sweep  
4        operation for visible configuration changes.

1        4.    The method of Claim 2, wherein said network includes  
2        a switch and said dynamically determining step includes  
3        the steps of:

4                detecting an addition of said component to a link of

5       said switch; and

6               in response to said detecting step, signaling said  
7       management system via a trap message that said component  
8       is on said network

1       5.    The method of Claim 2 further comprising the steps  
2       of associating said component to at least one partition  
3       and to an OS within said at least one partition.

1       6.    The method of Claim 5, further comprising the steps  
2       of:

3               detecting a partition of said network to which said  
4       component has been associated; and

5               notifying said OS of said component only when said  
6       OS is also associated with said partition.

1       7.    The method of Claim 6, further comprising the steps  
2       of:

3               updating a component table available to said OS with  
4       information about said component; and

5               providing OS support to all components registered in  
6       said component table.

1 8. A system for configuring a network, said system  
2 comprising:

3 a network manager that dynamically determines when a  
4 component is added to a node of said network and  
5 configures said network to provide support for said  
6 component; and

7 a network administration utility that, and in  
8 response to said network manager dynamically determining  
9 when a component is added, notifies an OS registered with  
10 said network administration utility that said component  
11 is added.

1 9. The system of Claim 8, wherein said network manager  
2 determines when said component is added by periodically  
3 monitoring said network for configuration changes.

1 10. The system of Claim 8, wherein said network manager  
2 determines when said component is added by receiving a  
3 packet from said component indicated that said component  
4 is present on said network.

1 11. The system of Claim 8, further comprising a  
2 registration table for registering said OS with said  
3 network administration utility for notification of an  
4 addition of a component; and

1 12. The system of Claim 11, further comprising:  
2 a partitioning mechanism that associates said  
3 component with one or more of a plurality of partitions  
4 of said network; and

5 wherein said network manager notifies said OS only

13. The system of Claim 12, further comprising a component registry available to said OS that is updated with information about said component when said component is detected, wherein said OS provides support to all components registered in said component registry to which it has access permission.

1 14. A network comprising:

2 a switch;

3 at least one node linked to said switch for adding  
4 components;

5 a network manager that dynamically determines when a  
6 component is added to said at least one node of said  
7 network and configures said network to provide support  
8 for said component;

9 an operating system; and

10 a network administration utility that, and in  
11 response to said network manager, dynamically determining  
12 when a component is added, notifies an OS registered with  
13 said network administration utility that said component  
14 is added.

15. The network of Claim 14, further comprising:

1 a partition agent that associates said component to  
2 one or more partitions of said network and controls  
3 access to said component via a partition monitoring  
4 function; and  
5

6 wherein said OS is notified of said component only  
7 when said OS has an access permission to a same one of  
8 said one or more partitions.

1 16. The network of Claim 14, wherein said network is a  
2 system area network.

1 17. A computer program product comprising:

2 a computer readable medium; and

3 program instructions on said computer readable  
4 medium for:

5 dynamically determining when a component is  
6 added to a node of a network; and

7 in response to said dynamically determining  
8 step, configuring said network to provide support  
9 for said component.

1 18. The computer program product of Claim 17, further  
2 comprising program instructions for:

3 registering an OS with an management system of said  
4 network to provide a notification when said component is  
5 added; and

6 alerting said OS that said component is added during  
7 said configuring step.

1 19. The computer program product of Claim 18 further  
2 comprising program instructions for associating said  
3 component to at least one partition and to an OS within  
4 said at least one partition.

1 20. The computer program product of Claim 19, further  
2 comprising program instructions for:

3 detecting a partition of said network to which said  
4 component has been associated; and



5           notifying said OS of said component only when said  
6           OS is also associated with said partition.

1           21. The computer program product of Claim 19, further  
2           comprising program instructions for:

3           detecting a partition of said network to which said  
4           component has been associated; and

5           notifying said OS only when said OS has access  
6           permission for said partition and said component.

1           22. The computer program product of Claim 21, further  
2           comprising program instructions for:

3           updating a component table available to said OS with  
4           information about said component; and

5           providing OS support to all components registered in  
said component table.

**ABSTRACT OF THE DISCLOSURE**  
**METHOD AND SYSTEM FOR INFORMING THE OPERATING SYSTEM IN A**  
**SYSTEM AREA NETWORK WHEN A NEW DEVICE IS CONNECTED**

5

A method and system are provided for dynamically informing an operating system (OS) of a distributed computer system, when a (new) device is added on the network. An OS registers with the subnet administrator (SA) to be notified when a new component is added to the subnet and placed in a partition to which the OS has access. The subnet manager (SM) initializes and configures the new component that is added and, through the partition manager, associates the component to one or more partitions of the subnet. The SA then notifies the OSs with access to the partition(s) to which the component is associated and that have a right to access the new component that the new component is available to the OS. Following, the OS updates a database of available components to include the new component and establishes communication with the new component.

10

15

20

1/8

EKS38216795US

AU39-2000-0620US1

Figure 1:






Figure 2:



1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11

3/8

Figure 3: Channel Adapter



1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17

4/8

Figure 4: Work and Completion Queue Processing





FIG. 5



4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27

Figure 6






Figure 7.


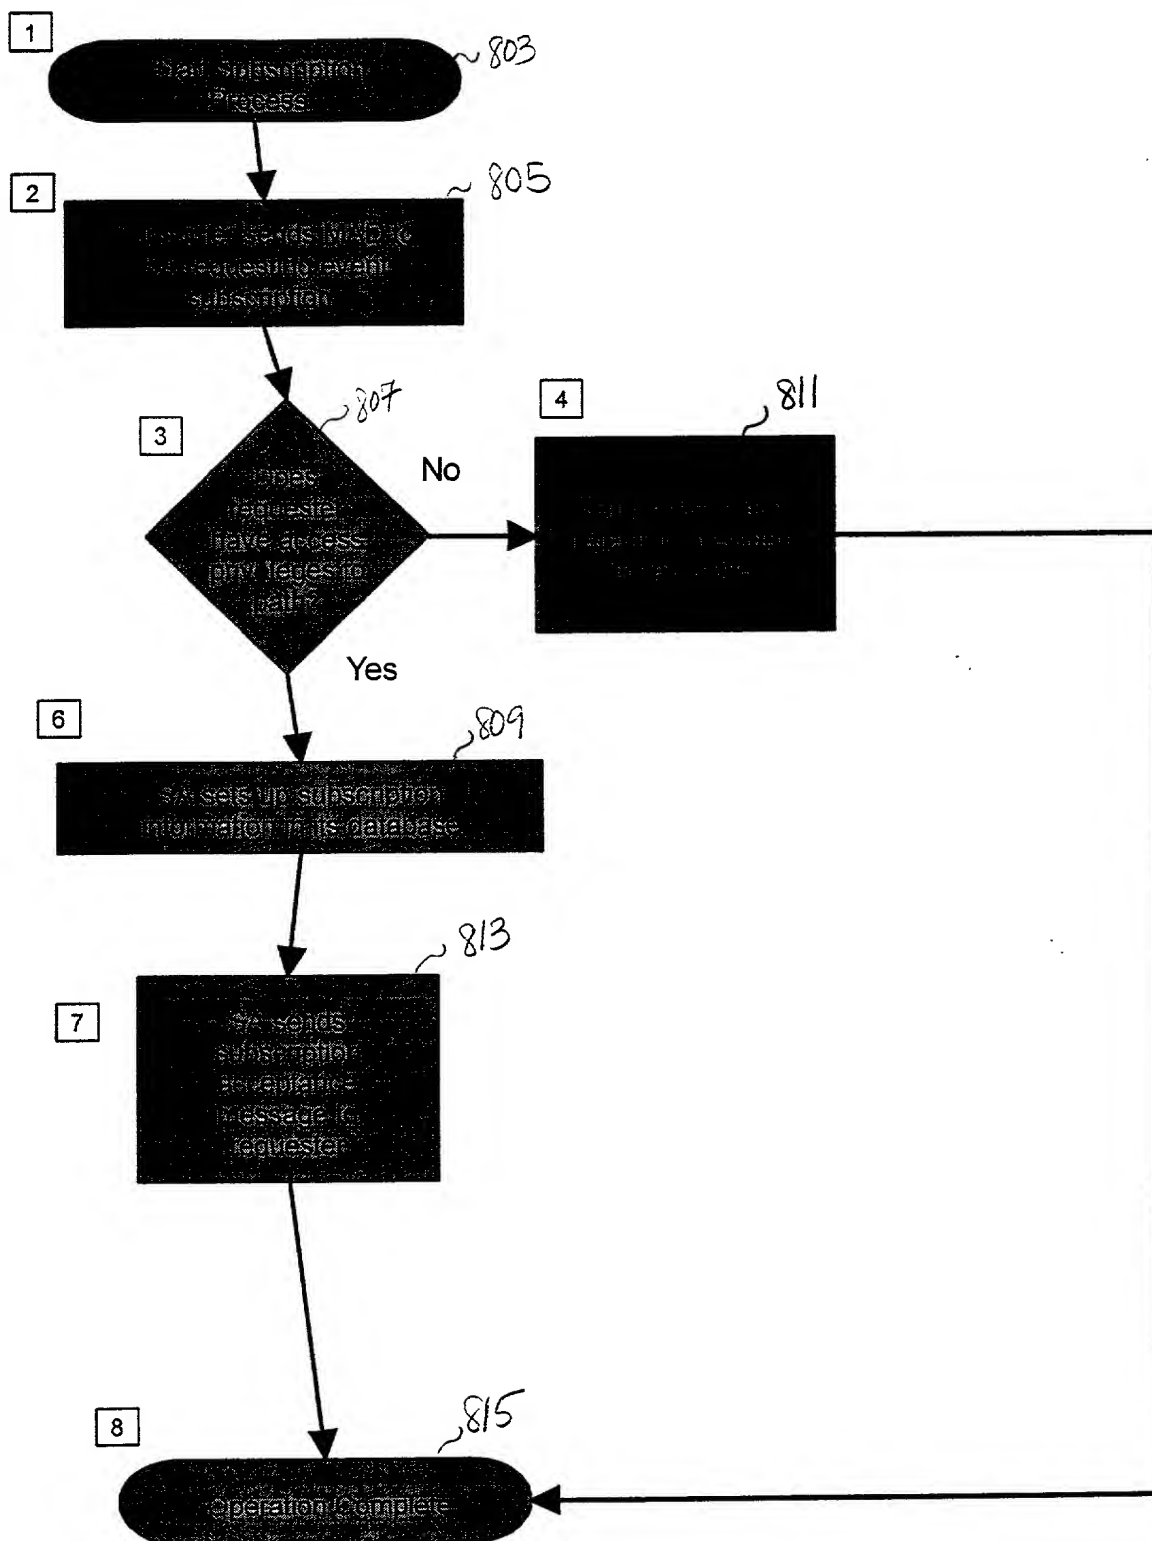


Figure 8



**DECLARATION AND POWER OF ATTORNEY FOR  
PATENT APPLICATION**

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name;

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled

**METHOD AND SYSTEM FOR INFORMING AN OPERATING SYSTEM IN A SYSTEM  
AREA NETWORK WHEN A NEW DEVICE IS CONNECTED**

the specification of which (check one)

X is attached hereto.

\_\_\_\_\_ was filed on \_\_\_\_\_  
as Application Serial No. \_\_\_\_\_  
and was amended on \_\_\_\_\_  
(if applicable)

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to the patentability of this application in accordance with Title 37, Code of Federal Regulations, §1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, §119 of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

Prior Foreign Application(s):

Priority Claimed

\_\_\_\_\_ Yes \_\_\_\_\_ No  
(Number) (Country) (Day/Month/Year)

I hereby claim the benefit under Title 35, United States Code, §120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, §112, I acknowledge the duty to disclose information material to

the patentability of this application as defined in Title 37, Code of Federal Regulations, §1.56 which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

(Application Serial #)

(Filing Date)

(Status)

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

POWER OF ATTORNEY: As a named inventor, I hereby appoint the following attorneys and/or agents to prosecute this application and transact all business in the Patent and Trademark Office connected therewith.

John W. Henderson, Jr., Reg. No. 26,907; Thomas E. Tyson, Reg. No. 28,543; Robert M. Carwell, Reg. No. 28,499; Jeffrey S. LaBaw, Reg. No. 31,633; Douglas H. Lefevre, Reg. No. 26,193; Casimer K. Salys, Reg. No. 28,900; David A. Mims, Jr., Reg. No. 32,708; Volel Emile, Reg. No. 39,969; James H. Barksdale, Jr. Reg. No. 24,091; Leslie A. Van Leeuwen, Reg. No. 42,196; Marilyn S. Dawkins, Reg. No. 31,140; Mark E. McBurney, Reg. No. 33,114; Christopher A. Hughes, Reg. No. 26,914; Edward A. Pennington, Reg. No. 32,588; John E. Hoel, Reg. No. 26,279; Joseph C. Redmond, Jr., Reg. No. 18,753; Andrew J. Dillon, Reg. No. 29,634; Daniel E. Venglarik, Reg. No. 39,409; Jack V. Musgrove, Reg. No. 31,986; Brian F. Russell, Reg. No. 40,796; Matthew W. Baca, Reg. No. 42,277; Antony P. Ng, Reg. No. 43,427; Michael R. Barre, Reg. No. 44,023; Andrew Mitchell Harris, Reg. No. 42,638; Richard McCain, Reg. No. 43,785; Michael E. Noe, Jr., Reg. No. 44,975.

Send correspondence to: Andrew J. Dillon, Esq., FELSMAN, BRADLEY, VADEN, GUNTER & DILLON, LLP, Suite 350 Lakewood on the Park, 7600B North Capital of Texas Highway, Austin, Texas 78731, and direct all telephone calls to Atty. Andrew J. Dillon, (512) 343-6116.

FULL NAME OF SOLE OR FIRST INVENTOR: GREGORY MICHAEL NORDSTROM

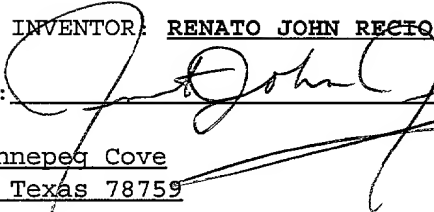
INVENTORS SIGNATURE: \_\_\_\_\_ DATE: \_\_\_\_\_

RESIDENCE: 1001 Golf Course Road  
Pine Island, Minnesota 55963

CITIZENSHIP: U.S.A.

POST OFFICE ADDRESS: 1001 Golf Course Road  
Pine Island, Minnesota 55963

FULL NAME OF SECOND INVENTOR: RENATO JOHN RECTO

INVENTORS SIGNATURE:  DATE: 10/16/00

RESIDENCE: 6707 Winnepeg Cove  
Austin, Texas 78759

CITIZENSHIP: U.S.A.

POST OFFICE ADDRESS: 6707 Winnepeg Cove  
Austin, Texas 78759

FULL NAME OF THIRD INVENTOR: GREGORY FRANCIS PFISTER

INVENTORS SIGNATURE:  DATE: 10/13/00

RESIDENCE: 5905 Sir Ivor Cove  
Austin, Texas 78746

CITIZENSHIP: U.S.A.

POST OFFICE ADDRESS: 5905 Sir Ivor Cove  
Austin, Texas 78746

FULL NAME OF FOURTH INVENTOR: ORVALLE THEODORE KIRBY

INVENTORS SIGNATURE: \_\_\_\_\_ DATE: \_\_\_\_\_

RESIDENCE: 9436 Macon Road  
Raleigh, North Carolina 27613

CITIZENSHIP: U.S.A.

POST OFFICE ADDRESS: 9436 Macon Road  
Raleigh, North Carolina 27613

FULL NAME OF FIFTH INVENTOR: STEVEN MARK THURBER

INVENTORS SIGNATURE:  DATE: 10/12/2000

RESIDENCE: 8308 Ephraim Road  
Austin, Texas 78727

CITIZENSHIP: U.S.A.

POST OFFICE ADDRESS: 8308 Ephraim Road  
Austin, Texas 78727

DOCKET NUMBER: AUS920000620US1

DECLARATION AND POWER OF ATTORNEY FOR  
PATENT APPLICATION

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name;

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled

METHOD AND SYSTEM FOR INFORMING AN OPERATING SYSTEM IN A SYSTEM  
AREA NETWORK WHEN A NEW DEVICE IS CONNECTED

the specification of which (check one)

X is attached hereto.

\_\_\_\_\_ was filed on \_\_\_\_\_  
as Application Serial No. \_\_\_\_\_  
and was amended on \_\_\_\_\_  
(if applicable)

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to the patentability of this application in accordance with Title 37, Code of Federal Regulations, §1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, §119 of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

Prior Foreign Application(s):			Priority Claimed
_____	_____	_____	Yes No
(Number)	(Country)	(Day/Month/Year)	

I hereby claim the benefit under Title 35, United States Code, §120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, §112, I acknowledge the duty to disclose information material to

DOCKET NUMBER: AUS920000620US1

the patentability of this application as defined in Title 37, Code of Federal Regulations, §1.56 which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

(Application Serial #) (Filing Date) (Status)

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

POWER OF ATTORNEY: As a named inventor, I hereby appoint the following attorneys and/or agents to prosecute this application and transact all business in the Patent and Trademark Office connected therewith.

John W. Henderson, Jr., Reg. No. 26,907; Thomas E. Tyson, Reg. No. 28,543; Robert M. Carwell, Reg. No. 28,499; Jeffrey S. LaBaw, Reg. No. 31,633; Douglas H. Lefevre, Reg. No. 26,193; Casimer K. Salys, Reg. No. 28,900; David A. Mims, Jr., Reg. No. 32,708; Volel Emile, Reg. No. 39,969; James H. Barksdale, Jr. Reg. No. 24,091; Leslie A. Van Leeuwen, Reg. No. 42,196; Marilyn S. Dawkins, Reg. No. 31,140; Mark E. McBurney, Reg. No. 33,114; Christopher A. Hughes, Reg. No. 26,914; Edward A. Pennington, Reg. No. 32,588; John E. Hoel, Reg. No. 26,279; Joseph C. Redmond, Jr., Reg. No. 18,753; Andrew J. Dillon, Reg. No. 29,634; Daniel E. Venglarik, Reg. No. 39,409; Jack V. Musgrove, Reg. No. 31,986; Brian F. Russell, Reg. No. 40,796; Matthew W. Baca, Reg. No. 42,277; Antony P. Ng, Reg. No. 43,427; Michael R. Barre, Reg. No. 44,023; Andrew Mitchell Harris, Reg. No. 42,638; Richard McCain, Reg. No. 43,785; Michael E. Noe, Jr., Reg. No. 44,975.

Send correspondence to: Andrew J. Dillon, Esq., FELSMAN, BRADLEY, VADEN, GUNTER & DILLON, LLP, Suite 350 Lakewood on the Park, 7600B North Capital of Texas Highway, Austin, Texas 78731, and direct all telephone calls to Atty. Andrew J. Dillon, (512) 343-6116.

FULL NAME OF SOLE OR FIRST INVENTOR: GREGORY MICHAEL NORDSTROM

INVENTORS SIGNATURE: Gregory Michael Nordstrom

DATE: October 17, 2000

RESIDENCE: 1001 Golf Course Road  
Pine Island, Minnesota 55963

CITIZENSHIP: U.S.A.

POST OFFICE ADDRESS: 1001 Golf Course Road  
Pine Island, Minnesota 55963

DOCKET NUMBER: AUS920000620US1

FULL NAME OF SECOND INVENTOR: RENATO JOHN RECIO

INVENTORS SIGNATURE: \_\_\_\_\_ DATE: \_\_\_\_\_

RESIDENCE: 6707 Winnepeg Cove  
Austin, Texas 78759

CITIZENSHIP: U.S.A.

POST OFFICE ADDRESS: 6707 Winnepeg Cove  
Austin, Texas 78759

FULL NAME OF THIRD INVENTOR: GREGORY FRANCIS PFISTER

INVENTORS SIGNATURE: \_\_\_\_\_ DATE: \_\_\_\_\_

RESIDENCE: 5905 Sir Ivor Cove  
Austin, Texas 78746

CITIZENSHIP: U.S.A.

POST OFFICE ADDRESS: 5905 Sir Ivor Cove  
Austin, Texas 78746

FULL NAME OF FOURTH INVENTOR: ORVILLE THEODORE KIRBY

INVENTORS SIGNATURE: \_\_\_\_\_ DATE: \_\_\_\_\_

RESIDENCE: 9436 Macon Road  
Raleigh, North Carolina 27613

CITIZENSHIP: U.S.A.

POST OFFICE ADDRESS: 9436 Macon Road  
Raleigh, North Carolina 27613

FULL NAME OF FIFTH INVENTOR: STEVEN MARK THURBER

INVENTORS SIGNATURE: \_\_\_\_\_ DATE: \_\_\_\_\_

RESIDENCE: 8308 Ephraim Road  
Austin, Texas 78727

CITIZENSHIP: U.S.A.

POST OFFICE ADDRESS: 8308 Ephraim Road  
Austin, Texas 78727



DOCKET NUMBER: AUS920000620US1

DECLARATION AND POWER OF ATTORNEY FOR  
PATENT APPLICATION

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name;

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled

METHOD AND SYSTEM FOR INFORMING AN OPERATING SYSTEM IN A SYSTEM  
AREA NETWORK WHEN A NEW DEVICE IS CONNECTED

the specification of which (check one)

X is attached hereto.

\_\_\_\_\_ was filed on \_\_\_\_\_  
as Application Serial No. \_\_\_\_\_  
and was amended on \_\_\_\_\_  
(if applicable)

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to the patentability of this application in accordance with Title 37, Code of Federal Regulations, §1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, §119 of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

Prior Foreign Application(s): \_\_\_\_\_ Priority Claimed

\_\_\_\_\_ Yes \_\_\_\_\_ No  
(Number) (Country) (Day/Month/Year)

I hereby claim the benefit under Title 35, United States Code, §120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, §112, I acknowledge the duty to disclose information material to

DOCKET NUMBER: AUS920000620US1

the patentability of this application as defined in Title 37, Code of Federal Regulations, §1.56 which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

(Application Serial #)	(Filing Date)	(Status)
------------------------	---------------	----------

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

POWER OF ATTORNEY: As a named inventor, I hereby appoint the following attorneys and/or agents to prosecute this application and transact all business in the Patent and Trademark Office connected therewith.

John W. Henderson, Jr., Reg. No. 26,907; Thomas E. Tyson, Reg. No. 28,543; Robert M. Carwell, Reg. No. 28,499; Jeffrey S. LaBaw, Reg. No. 31,633; Douglas H. Lefevre, Reg. No. 26,193; Casimer K. Salys, Reg. No. 28,900; David A. Mims, Jr., Reg. No. 32,708; Volel Emile, Reg. No. 39,969; James H. Barksdale, Jr. Reg. No. 24,091; Leslie A. Van Leeuwen, Reg. No. 42,196; Marilyn S. Dawkins, Reg. No. 31,140; Mark E. McBurney, Reg. No. 33,114; Christopher A. Hughes, Reg. No. 26,914; Edward A. Pennington, Reg. No. 32,588; John E. Hoel, Reg. No. 26,279; Joseph C. Redmond, Jr., Reg. No. 18,753; Andrew J. Dillon, Reg. No. 29,634; Daniel E. Venglarik, Reg. No. 39,409; Jack V. Musgrove, Reg. No. 31,986; Brian F. Russell, Reg. No. 40,796; Matthew W. Baca, Reg. No. 42,277; Antony P. Ng, Reg. No. 43,427; Michael R. Barre, Reg. No. 44,023; Andrew Mitchell Harris, Reg. No. 42,638; Richard McCain, Reg. No. 43,785; Michael E. Noe, Jr., Reg. No. 44,975.

Send correspondence to: Andrew J. Dillon, Esq., FELSMAN, BRADLEY, VADEN, GUNTER & DILLON, LLP, Suite 350 Lakewood on the Park, 7600E North Capital of Texas Highway, Austin, Texas 78731, and direct all telephone calls to Atty. Andrew J. Dillon, (512) 343-6116.

FULL NAME OF SOLE OR FIRST INVENTOR: GREGORY MICHAEL NORDSTROM

INVENTORS SIGNATURE: \_\_\_\_\_ DATE: \_\_\_\_\_

RESIDENCE: 1001 Golf Course Road  
Pine Island, Minnesota 55963

CITIZENSHIP: U.S.A.

POST OFFICE ADDRESS: 1001 Golf Course Road  
Pine Island, Minnesota 55963

DOCKET NUMBER: AUS920000620US1

FULL NAME OF SECOND INVENTOR: RENATO JOHN RECIO

INVENTORS SIGNATURE: \_\_\_\_\_ DATE: \_\_\_\_\_

RESIDENCE: 6707 Winnepeg Cove  
Austin, Texas 78759

CITIZENSHIP: U.S.A.

POST OFFICE ADDRESS: 6707 Winnepeg Cove  
Austin, Texas 78759

FULL NAME OF THIRD INVENTOR: GREGORY FRANCIS PFISTER

INVENTORS SIGNATURE: \_\_\_\_\_ DATE: \_\_\_\_\_

RESIDENCE: 5905 Sir Ivor Cove  
Austin, Texas 78746

CITIZENSHIP: U.S.A.

POST OFFICE ADDRESS: 5905 Sir Ivor Cove  
Austin, Texas 78746

FULL NAME OF FOURTH INVENTOR: ORVALLE THEODORE KIRBY

INVENTORS SIGNATURE: Orville Hudson/Gil DATE: October 13, 2020

RESIDENCE: 9436 Macon Road  
Raleigh, North Carolina 27613

CITIZENSHIP: U.S.A.

POST OFFICE ADDRESS: 9436 Macon Road  
Raleigh, North Carolina 27613

FULL NAME OF FIFTH INVENTOR: STEVEN MARK THURBER

INVENTORS SIGNATURE: \_\_\_\_\_ DATE: \_\_\_\_\_

RESIDENCE: 8308 Ephraim Road  
Austin, Texas 78727

CITIZENSHIP: U.S.A.

POST OFFICE ADDRESS: 8308 Ephraim Road  
Austin, Texas 78727